



Deposited via The University of Leeds.

White Rose Research Online URL for this paper:

<https://eprints.whiterose.ac.uk/id/eprint/2334/>

Monograph:

Bonsall, P.W., Kirby, H.R. and Kwan, R.S. (1986) Applications of Expert Systems in Transport. Working Paper. Institute of Transport Studies, University of Leeds , Leeds, UK.

Working Paper 216

Reuse

See Attached

Takedown

If you consider content in White Rose Research Online to be in breach of UK law, please notify us by emailing eprints@whiterose.ac.uk including the URL of the record and the reason for the withdrawal request.



White Rose Research Online

<http://eprints.whiterose.ac.uk/>

ITS

[Institute of Transport Studies](#)

University of Leeds

This is an ITS Working Paper produced and published by the University of Leeds. ITS Working Papers are intended to provide information and encourage discussion on a topic in advance of formal publication. They represent only the views of the authors, and do not necessarily reflect the views or approval of the sponsors.

White Rose Repository URL for this paper:

<http://eprints.whiterose.ac.uk/2334/>

Published paper

Bonsall, P.W., Kirby, H.R., Kwan, R.S. (1986) *Applications of Expert Systems in Transport*. Institute of Transport Studies, University of Leeds. Working Paper 216

Working Paper 216

April 1986

APPLICATIONS OF EXPERT SYSTEMS IN TRANSPORT

P.W. BONSTALL, H.R. KIRBY and R.S. KWAN

This paper supersedes ITS Technical Note 145 (by MD Wheatley) and Technical Note 164 (by PW Bonsall); and comprises a pre-publication version of a paper entitled

"The role of expert systems in transport"

by PW Bonsall and HR Kirby (to appear in Information Technology Applications in Transport by VNU Science Press, 1986, edited by PW Bonsall and M Bell); together with an Appendix by RS Kwan.

If cited, reference should be given to the published version whose details are above.

ITS Working Papers are intended to provide information and encourage discussion on a topic in advance of formal publication. They represent only the views of the authors and do not necessarily reflect the views or approval of sponsors.

This work was sponsored by the University of Leeds and the Science and Engineering Research Council

THE ROLE OF EXPERT SYSTEMS IN TRANSPORT

P W Bonsall

H R Kirby

Institute for Transport Studies
University of Leeds

To be published in Information Technology Applications in Transport,
VNU, Science Press, 1986. Edited by P W Bonsall and M Bell.

INTRODUCTION

Background

Experienced judgement and specialist knowledge are essential to the proper specification, understanding and interpretation of data and computer analyses. The human expert has traditionally supplied this knowledge and judgement with the computer doing the necessary number-crunching. However, artificial intelligence (AI) research provides ways of embodying this knowledge and judgement within computer programs. Despite an early lead in the field, UK research and development into AI techniques was held back in the 1970s when the then Science Research Council took the view that the 'combinatorial explosion' of possibilities would be an insurmountable obstacle to AI development. But in America and Japan research continued, and the surge of interest in the 1980s has been a consequence of the 'Fifth Generation Computer' research programme initiated by Japan (Feigenbaum and McCorduck, 1984). This led in Europe to the ESPRIT programme of advanced technology research, and in the UK to the Alvey programme (Department of Industry, 1982). As a result, all sectors of industry have been encouraged to consider how such advanced technology can be applied, and the transport industry is no exception.

This paper sets out to explain some of the relevant techniques in simple terms, and to describe a number of situations in which transport planning and operations might be helped through their use, illustrating this by reference to the pioneering work going on in transport applications in the USA, Britain and Australia.

What the phrases mean.

Artificial intelligence implies intelligent-seeming behaviour in computers. The scope is clearly vast, including such subjects as image recognition, the understanding of natural language and the control of robots.

AI research has spawned a number of specialist languages and tools for different classes of problem. The tools most likely to be useful for transport planning and operations are known as expert systems. They involve rules and relationships devised to express human knowledge and judgemental processes. Because

they involve the application of human knowledge to a problem, expert systems are often more modestly referred to as knowledge-based systems (KBS). The phrase 'knowledge-based systems' is in principle broader than that of expert systems, since it includes those situations in which the knowledge is readily available, say in a manual or a set of regulations, rather than in the experience of an expert. Another phrase is 'Intelligent Knowledge-Based Systems' (IKBS), which has become the new vogue term following its adoption in the Alvey Report (Department of Industry, 1982). IKBS was there defined to be "a system which uses inference to apply knowledge to perform a task"; it thus implies advanced kinds of expert system.

WHAT CAN EXPERT SYSTEMS DO?

Some examples.

Different kinds of knowledge-based system have evolved over the years, with differing characteristics. Some idea of the capabilities of the systems is apparent from the following list of types of problem to which they have been addressed. (The names of some of the pioneering systems are given in parentheses; these are described in most AI text-books).

- Diagnosis of disease (MYCIN); of electrical circuits (EL).
- Enumeration of molecular structure (DENDRAL).
- Detection of minerals (PROSPECTOR).
- Configuration of computer systems (XCON).

The new tools and techniques imply a fresh perspective on what is possible with computers by clients, users and programmers alike.

Client's and users' perspectives.

Knowledge acquisition and dissemination. The main advantage of an expert system for a client is that it provides a way of disseminating somebody's experienced knowledge or judgment concerning a problem to less experienced staff, or of bringing together the experience of several different experts to bear on a complex problem. This will not only make for better and speedier decision-making by the user, but also releases the expert's time for other tasks - or ensures that the expert's skills can be retained in the organisation when she/he leaves. The rules of thumb are most obviously derived directly from the expert(s); but this process can be extremely difficult and time-consuming and, in some circumstances, speedy system development can be facilitated by recording the decisions made by the expert in a given set of circumstances and then inducing the rules he is following from an analysis of his actions.

Even before the expert system itself is constructed, an advantage may be gained by the organisation if the expert's rules of thumb and judgmental values, which are usually implicit and may not at first be recognised as such by him, are made explicit by the knowledge acquisition process. Because they encapsulate expertise, expert systems also have a useful role in training less experienced staff.

Updating and maintenance. Since the knowledge of the expert(s) and the way in which it all fits together usually becomes apparent only gradually, and indeed will change over time in the light of experience and exogenous factors, an important feature of the expert system from the client's point of view is that it is easy to update the knowledge base. Knowledge-based systems are so constructed that the rules and relationships in them are transparent, and readily modified.

Treatment of uncertainty. In real life, input data is often imprecise either because of measurement error or because of fluctuations in the basic data (eg one might estimate traffic flow at 1500 vehicles/hour but it could be between 1400 and 1600 vehicles/h). Sometimes the data is more naturally described qualitatively (eg a 'high' flow) rather than quantitatively. Also, for given values of input data (no matter how precise), there may be several possible courses of action (or outcomes). Expert systems can be set up to cope with both uncertainty in input data and with uncertainty of outcome in a statistically appropriate way. (For a discussion of how this is done see below.)

User-friendliness. A prime desideratum of any computer system is that it should be easy for the client's staff to use. Clearly, this is particularly important if the intention is to make the system available to non-specialists. An expert system is not in itself user-friendly; but it can contribute to user-friendliness by ensuring that, in the machine's dialogue with the user, it asks only the relevant questions, is forgiving in its treatment of mistakes, perceives inconsistencies in the user's answers to questions, and checks for gaps in its own knowledge. It can also explain and justify its reasoning on request. A good expert system will have the user thinking it was clever enough to think of something he/she had not thought of, but recognises as true once seen.

Programmer's perspective.

Languages. Programmers familiar with conventional computing for engineering and transport applications will find that expert systems require a rather different style of approach. Traditionally they will have been concerned with the manipulation of numbers, and will probably have used the Fortran (or, more recently, Pascal) language for this purpose. These languages are fine for engineering or scientific calculations; but they are rather clumsy at manipulating rules or relationships. The AI community have developed other computer languages that are much better able to do this. One is LISP, which was originally designed (in 1958) for processing lists. Modern-day implementations, available through such products as LOOPS, KNOWLEDGE CRAFT, and KEE, are very different, in that they build powerful tools on top of LISP. They offer much more powerful and easy-to-use ways of representing knowledge, as well as good project-support environments including 'menus' and 'windows'. Another main language is PROLOG, which is designed for processing logical relationships; this is the

language adopted in the Japanese research programme for a Fifth Generation Computer. POPLOG provides an environment within which codes from either of these two languages (or another - POP11) can be tied together. This environment provides editing facilities and enables small chunks of code to be edited, compiled and tested independently. SMALLTALK-80 is an example of a class of so-called 'object-oriented' languages which are designed to facilitate problem description by, for example, allowing a newly defined object to 'inherit' some of the characteristics of previously defined objects.

System components. The main conceptual development involved is the idea that, like data, the rules are separated from the rest of the program. The collection of rules is known as the knowledge base; these fit into what is known as a shell. This separation of the knowledge (which may be represented in a variety of ways) from the mechanism for manipulating the knowledge makes it very much easier to introduce new rules, and to see the structural relationships involved in the knowledge base, compared say with a Fortran program to do the same job. The other main feature of an expert system is the mechanism involved in the manipulation of knowledge and drawing of inferences. Technically, this part of the expert system is known as the inference engine; some of its characteristics are described below.

Uncertainty in input. Expert systems can deal with imprecision in input data either by requiring the user to specify probability weights for different states of the input variables or, more rarely, by the use of sophisticated procedures based on fuzzy logic.

Uncertainty of outcome. While some inference mechanisms are deterministic involving what are known as 'production rules' (of the form 'IF A is true THEN B is true'), most mechanisms involve statistical procedures in order to make inferences. Thus for example the rule might be 'IF A is true THEN B will be true with probability 0.7 AND C will be true with probability 0.3'. Although there are some alternative techniques available, each with their own advocates (see Mandani et al, 1985), most expert systems use the well-known statistical principle of Bayes' Theorem to manipulate the probabilities in the rule set so as to reach an overall judgement of the probability of a particular objective being true. This consistency and the appropriate weighting of probabilities is an important attribute of such systems.

Explanation and description. An important attribute of an expert system is that it can be asked to explain its reasoning. Since it is desirable, for ease of comprehension, that its rules and explanations are readily recognised, the computer has an enhanced descriptive role. The programmer will therefore have to spend a greater proportion of his/her effort on ensuring that the descriptive text and rules are readily understood, and that the system is developed in a user-friendly way.

THE SCOPE FOR APPLYING EXPERT SYSTEMS IN TRANSPORT

Context.

A number of authors (eg Wigan, 1983; Logie and Neffendorf 1984) have drawn attention to the possibility of applying expert systems to aid the study of transport problems. The work reported here is the outcome of a detailed examination of the scope for such applications, carried out at the Institute for Transport Studies (ITS) of the University of Leeds. The range of applications considered, though wide, naturally reflects the particular interests of the Institute; others considering a breadth of transport applications include Wigan (1985) at the Australian Road Research Board; Hendrickson et al (1985) at Carnegie-Mellon University, Philadelphia; and Yeh et al (1985) at the University of Washington, Seattle. Many other institutions are now responding to the challenge of adopting AI techniques for use in transport applications, particularly in the USA; this is indicated by the way in which the topic bubbles up in the papers reviewing new transportation research opportunities (Boyce, 1985). References to other applications are included below.

The following list summarises potential applications of expert systems considered in the ITS project. It should be noted that the list is arranged by type of application and therefore cuts across the technical distinctions between different types of expert system.

Enquiry Systems

- advice on regulations
- travel enquiries
- route advice
- advice on data sources

Analytical Advice and Interpretation Systems

- procedural and methodological advice
- interpretation

Design Systems

- infrastructure design
- network/junction design
- schedule design
- questionnaire design

Diagnostic/Prescriptive Systems

- road safety systems
- road maintenance
- structures and equipment maintenance

Identification Systems

- inventory and condition logging
- automatic identification

Control Systems

- traffic monitoring and control
- air traffic control
- survey control

Policy Support Systems

- multi-criteria decision making
- treatment of uncertainty
- consistency of policy

We now consider what might be involved in each of these applications. Because of the range of subjects covered, the

different applications are discussed in different levels of detail, but it is hoped this will stimulate further debate.

Enquiry Systems.

The 'expertise' inherent in a successful enquiry system is the ability to respond to an enquiry so as to provide the necessary information as quickly as possible. This implies an ability to tease out of the enquirer precisely what it is that he wishes to know, to access the relevant information and to pass it on clearly and succinctly. Of these the most difficult is the interpretation of the initial enquiry and the decision as to what information would be useful. In order to provide such a service it is of course necessary to have a sound and up-to-date knowledge of the field in question.

Expert systems have much to contribute here because of their ability to arrange their knowledge bases such that they 'know' the extent of their knowledge (a concept known as meta-knowledge); because their knowledge can be rapidly updated from a variety of sources; because they can respond to questions posed in a variety of ways; because they can respond flexibly (eg changing direction in response to additional or amended information from the enquirer) and because they can learn from experience (what is found useful in one session may be useful in similar circumstances in a subsequent session).

Another feature of expert systems that might commend their use in enquiry systems is their ability to act through a natural language interface whether typed or even spoken. Travel enquiries have often featured in trials of dialogue interpretation programs eg GUS (described by Bobrow et al, 1977) but in practice the complexities inherent in such systems may rule them out and leave menu selection as the most efficient mode of dialogue (see Rich, 1983 p 334).

Enquiry systems can of course be designed either as a complete service for lay enquirers, as training aids for novice clerks or as reference aids for more-or-less experienced clerks. In the latter case it is clearly important that the expert system is designed to provide accurate information about complex or infrequently requested topics more quickly than would be possible via more conventional means. Experience with enquiry systems based on conventional computing languages and data-base management systems has shown that, despite the benefits of having the necessary information organised in a database, the updating and accessing of the system may be so frustratingly tedious as to lead to such systems being ignored. There are, of course, several types of enquiry system - we mention only four. They range from simple advice systems that can be developed in a self-contained shell, to those that inter-relate with databases.

Advice on regulations. An expert system whose knowledge base contained regulations (eg building or planning regulations, vehicle construction-and-use regulations, customs regulations) would be relatively easy to set up since the knowledge is already codified. In its basic form, the main advantages of such a system over a physical book of rules are the ease with which it can be updated to include new or amended regulations, and the explanations it can

provide on its recommendations. A more ambitious system might incorporate case study and precedence information and might build up its own case law (seeking the important circumstances and antecedent regulations from its informants). Such a system might then be able to quote precedence as well as regulatory justification for its advice in given circumstances.

Travel enquiries. The intelligent front-end to a travel enquiry system ought to include the following features:

- optimum organisation of information so as to speed up the enquiry (eg with the system learning, from its experience with previous enquiries, that before holiday periods a high proportion of enquirers are likely to be interested in discount fares);
- ability to deal with requests posed in a variety of ways (eg 'I want to leave after x', 'I want to arrive by y', 'I want to travel in daylight' ...);
- ability to deal with fuzzy requests (eg 'what trains travel before that?');
- ability to devise sensible trade-offs (eg 'that train will get you in at 11 am', 'if you set off 5 minutes earlier you could get in at 10am');
- ability to explain in an alternative format any instructions not immediately understood (eg to give times in terms of the 12-hour rather than 24-hour clock);
- ability to provide subsidiary information if requested (eg whether changes of platform are likely to be required).

An additional benefit of expert systems in travel enquiries would be an ability to combine data from a variety of sources (eg timetables in a variety of formats from a variety of operators, temporary timetable changes, special discount fares, real time information on delayed services etc).

Our own experience at the Institute for Transport Studies in designing a simple railway timetable and fares enquiry system (TALK: Train Advice Leeds-Kings Cross) using a SAGE shell on a UNIX VAX computer convinced us of the limitations of the shell's facilities (see Kwan, 1985) and led us to propose a more ambitious system to be developed via the Alvey Directorate's IKBS Awareness Programme (Software Sciences et al, 1985). As a result, the TRACE Community Club has been formed to produce prototype enquiry systems not only for public transport services but also for holiday and travel sales.

Route advice. A number of conventionally constructed route advice systems already exist, some for private vehicles (eg those provided by automobile associations for their members) and some for public transport users (eg Williamson and Miller, 1981). These are based on minimum 'cost' path-finding algorithms and thus an amendment (eg to take account of temporary changes or interruptions to the network) is relatively expensive to incorporate. An expert system adviser might overcome this problem by amending only those routes which it knows to be affected by the change. Other desirable features might include an ability to define routes according to a variety of user-specified criteria (eg time, distance, reliability, avoidance of low bridges, etc); to define routes in the user's preferred terminology (eg 'take signs for Oxford' or '3rd left, then 2nd right...'); to provide subsidiary information if requested (eg

places of interest en route); to provide trade-off information (eg 'the route that best meets your instructions is ... however, if you are prepared to go through Basingstoke you could save 15 minutes') and to learn from experience (eg that people do not want to go through Basingstoke no matter how much time it saves!). Most of the features can, of course, be provided via conventional software (see for example Wootton and Brett's (1986) description of the UK Automobile Association's new system) and it remains to be seen whether the potential benefits of an expert system version (notably easier updating and learning from experience) are worthwhile.

Routing advice for multi-drop tours is discussed below under 'schedule design'.

Advice on data sources. Users of transport statistics have focussed attention on the widespread difficulty they have in becoming aware of the existence, relevance and availability of appropriate secondary data for a specific purpose (McLean, 1983). Although keyword indexing and computerised data bases have been of significant importance they clearly leave something to be desired in their habit of unearthing either 'too much' or 'not enough'. An expert system might be able to mimic the ability of a knowledgeable, efficient (and patient!) statistical sources expert or librarian to tease out of the enquirers precisely what they wish to know and to point them to the relevant data. A proposal to prototype such a system is currently under development at the Institute for Transport Studies.

Analytical Advice and Interpretation Systems.

Advice on technical procedures/methodologies, covering such matters as choice of appropriate survey techniques, statistical tests or types of model, would obviously rely on precedence and established procedures. At their simplest, such systems might appear little more than a computerised checklist of techniques or methodologies that have been regarded as suitable in similar circumstances. However, the main advantages of a more advanced expert system would be that it might probe for the determining circumstances and, in learning mode, might learn from the user's reactions to its proffered advice. Also, the system could, on the basis of accumulated experience and practice, take account of tradeoffs between any extra accuracy of more ambitious procedures and the extra cost and data requirements which might be incurred. Clearly, the advice system could go further than simply suggesting which techniques should be used; it could also give advice, if asked, on how they should be used. For example, a session with a sampling strategy advice system might begin with the user making statements on survey budget, desired confidence levels and anything known on the shape of the distributions to be sampled. The system might then ask for additional information (eg on network configuration) before suggesting the numbers of counts or interviews of various types to be carried at specified locations at specified times. The user might then respond on the feasibility of the suggested strategy (from which responses the system could learn for future reference).

Interpretation systems. An expert system might be designed not only to provide advice but, through its interface with the analytical procedures in question, might actually implement the advised procedures. Thus a system advising on the choice of statistical tests or analytical models might be designed to analyse the results of such models and hence to suggest a series of tests or models each based on the results of the previous one. (Thus, in the transport demand field, a statistical test on mode choice data might suggest that a lexicographic mode choice model be run in preference to a compensatory model). Particular attractions of such a system are that it would ensure consistency between models and the parameters/coefficients which they share and would clearly save the analyst considerable hours of work poring over voluminous computer output.

It may seem a logical progression to move from an advice system such as those outlined above to a system which might, without pausing for human intervention, requisition the data, carry out the tests and implement the models which it thinks appropriate. Whether this is a desirable goal must depend on whether the 'rules' which govern such actions can be specified. A halfway house would be an interpretation system whose task was just to assist the human in his assessment of information in data-bases and in model outputs (drawing his attention to the most significant features) and which, if designed to accept data in a variety of formats, could simplify the provision of information required by otherwise free standing analytical procedures.

Design Systems.

The design process is a complex activity, meaning different things in different contexts. The outcome is eventually an arrangement of objects to perform a given task and meeting certain criteria. Computer-aided design has hitherto embraced just the drafting aspects of this, concentrating on the use of interactive computer graphics. There is considerable scope for extending the computer's role via AI techniques to cover the selection and examination of alternative designs that meet the appropriate criteria (see Mostow, 1985).

Three features of expert systems that might prove particularly valuable in analytical design systems are:

- an intelligent front end to established design procedures that could deal with information presented or required in a variety of formats;
- a readily modified rule base that could limit the range of designs to be evaluated (eg through specific safety or capacity requirements or through constraints on project budget or duration); and
- an intelligent interpreter of the output of design exercises that could summarise and evaluate the important strengths and weaknesses of particular designs.

We now consider how four different design systems might benefit from the application of expert systems technology.

Infrastructure design systems. The knowledge base for designing infrastructure would include information on the components, how they relate to each other, what overall criteria should be met and

what measures of performance should be calculated. Examples include bridge design (for a pilot expert system on which see Welch and Biswas, 1986), car-park layout, and loading-bay configurations. System features might include:

- user specification of major parameters such as cost, capacity, maximum overall dimensions, site configuration;
- user specification of special features and constraints (eg that it must be constructable within 6 weeks, that there is a supply of cheap aggregate available nearby);
- draft designs suggested by the system (in learning mode constructive user comment on the draft designs would ensure enhancement of their standard in future applications);
- user selection of preferred design;
- system requests for further information as necessary;
- production of final design.

Engineering calculations (eg of stress for a bridge, or of curvature for a road) would need to be embodied within the overall design process as would user-friendly interactive graphics features.

Network/junction design systems. Such systems would need to consider not only the dimensions of space and time but would also have to reflect the fact that facility usage may be dependent on the design (eg traffic flows being affected by road network design). Interfacing with traffic/travel demand models and data will therefore be an important feature of the overall system.

An example of the use of an expert system as a design aid is the experimental program TRALI, developed at Carnegie-Mellon University. TRALI serves as a 'traffic signalling assistant' by suggesting sequences of phases for an isolated traffic signal (Zozaya-Gorostiza and Hendrickson, 1985); its expert system features are mainly in respect of rule-based limits on the range of permitted designs. Further work at Leeds is expected to development this approach; the main advantage is that flexibility is facilitated by having the rules in a readily modified rule-base rather than subsumed in the code of more conventional traffic signal optimising, and sequence design programs.

Schedule design. Routing for a multi-drop delivery is of course the 'travelling salesman' class of problem. For practitioners the problem is compounded with that of scheduling, because different vehicles in the fleet have different availabilities and capacities. In practice, schedulers are likely to rely on empirical rules of thumb rather than on algorithms, as these have hitherto been rather clumsy and time-consuming to use - especially if constraints are imposed on the time at which drops may be made. A knowledge-based approach, probably coupled with algorithms, may provide a useful way of harnessing an experienced scheduler's skills and making them more widely available in an organisation, if implemented on a micro-computer. Golden and Baker (in Boyce, 1985, p 407) and Polak and Jessop (1986) have also drawn attention to the potential role of expert systems in logistics.

The scheduling of buses and bus crews raises similar issues. In both the public transport and the freight situation, the

knowledge base would need to incorporate experience from previous practice, and knowledge of local working practices and union agreements.

Questionnaire design systems. We have considered whether there might be a role for expert systems in questionnaire design and conclude that, to be useful, they would require the following features:

- user specification of the items of data sought;
- system suggestion of draft questionnaire wording (drawn from a menu of possible questions stored in its knowledge base): the system would previously have checked which items are to be asked of which respondents in which circumstances;
- in learning mode, the system might invite user comments on the format of the proposed questions and their relationship one to another: any comments being incorporated into the system's knowledge base. The system might, if it came across a data item it had not yet met, or if its advice on wording were criticised, ask the user to provide a suggested wording;
- a substantial part of the system would be concerned with ensuring that the component questions fit properly together: user feedback on this would again allow the system to build on its own experience.

However, unless one were designing a system with highly specific objectives (eg how to ask a question about income), it seems most unlikely that a system could satisfactorily cope with the typically very varied nature of questionnaires and subject matter. A system based in conventional computing, such as QUESTMAST (Centre for Educational Sociology, 1984), may well be a sufficiently helpful design aid for survey documentation, coding manuals, variable definitions and specification of statistical analyses.

Diagnostic/Prescriptive Systems.

The scope for simple expert systems to advise on issues such as road maintenance treatment and environmental assessment has been indicated by several authors (see for example, Wigan 1982 and 1985; Hendrickson et al, 1985; Harris et al, 1985). Appropriate features would be:

- an ability to assimilate data from a variety of sources, to diagnose problems, preferably in advance of actual failure, and to suggest cost effective solutions to avoid any (further) failure;
- calculation of the probabilities of failure/accident if given remedial measures are (or are not) taken;
- a learning mode wherein human experts feed in rules of thumb, case histories and other information;
- a knowledge updating mode wherein the system monitors the performance of its own prescribed solutions.

Examples include:

Road safety systems. The aim here would be to diagnose potential accident blackspots and hazards and prescribe remedial measures. Data input might include accident case histories and site characteristics (infrastructure, traffic, weather). The site characteristics data might well be continually updated from on-line databases. An investigation of such systems is currently underway at University College London.

Road maintenance systems Data inputs for these might include structural history (age, materials used, strata thickness, ground conditions, previous maintenance) cumulative traffic flow (by weight), weather conditions and (in a sample of locations) embedded instrumentation to provide continuous data on porosity, compaction, deformation etc. Moavenzadeh (in Boyce, 1985, p 505).

The integration of instrumentation with on-line monitoring through the medium of expert systems has been suggested by pilot work on a small-scale adviser is being conducted in the Department of Civil Engineering at the University of Leeds.

Structures and equipment maintenance systems. Data inputs for these would be analogous to those described for road maintenance systems. Clearly systems with this remit would normally be produced from within the structural and mechanical engineering fraternity. Nevertheless, there are several examples of structures and equipment which are particularly oriented to transport (eg bridges, rail track, vehicles, etc) and which might warrant specialised systems. Wood (1985), for example, has discussed an expert system for bus maintenance; the US Army have developed an expert system for advising on railway track maintenance.

Identification Systems.

Inventory and condition logging systems might provide an aid to relatively unskilled enumerators engaged in inventory and condition surveys (or even highly skilled enumerators engaged in highly complex inventory and condition surveys). They would enable the enumerator to describe the inventory item in his own terms and the system would then prompt him with appropriate questions to determine exactly what was being described and what its condition was.

The essential features of the system would be:

- acceptance of a variety of imprecise definitions in free text;
- determination of the most significant questions so as to home in on the precise nature of the item being described (be it height, colour, serial number etc);
- ability to be used on site or in the field;
- a learning mode within which the system could learn new terms (eg lamp standard = lamp post) and learn to accept imprecision where precision is unnecessary (a 'yellow' light might not need to be distinguished from an 'orange' one for some purposes);
- easy updating of the knowledge base to include items newly introduced (eg a new type of road sign) or not yet met (eg an obsolete type of traffic light), or to allow for new knowledge (eg that concrete spalling is preceded by such-and-such a symptom).

The advantages of such a system over a decision-tree based algorithm would be its flexibility and efficiency, but again, it is not clear whether these advantages would justify the extra effort required.

Automatic identification systems such as image processing, voice recognition or recognition of patterns from inductive loops are currently being developed to aid the identification of classes of vehicle, items of inventory and even individual vehicles. They

could each be enhanced by being linked with expert systems. The linkage between pattern recognition and IKBS is a contentious issue of fundamental research well beyond the scope of this paper, but even without recourse to such techniques it is clear that an expert system could help resolve ambiguities in the output of an automatic identification system. In essence such aid would amount to the system saying "in the circumstances it is probably x". The circumstances in question could cover whether:

- anything is known about the population being observed (eg at a previous observation point a given mix of vehicles might have been observed);

- anything expected/unexpected in the population (eg one might not expect to see a fire engine in a multi-storey car park or a fire hydrant on a motorway).

The training of such systems could obviously start by incorporation of rules of thumb (eg local registration plates predominate during peak hours) and would be an on-going process whereby the data base of contextual information linked with successful identifications is continuously updated.

Control Systems.

This group of applications are characterised by their use of on-line data to optimise the performance of a system. They would derive their knowledge of how to achieve that optimisation from encoded formulae and relationships supplemented by case histories including those to which they have input advice.

Traffic monitoring and control is an obvious example of such an application. Most of the 'advice' would doubtless best be effected automatically (eg rephasing of signals in a UTC system) but it might be desirable sometimes to issue advice that could only be effected by human intervention (eg "check a particular stretch of road for physical obstruction"). The data input to the expert system would be most obviously via automatic sensors (eg flow and queue detectors) but one can imagine an expert system which was able to request extra items of data from roving human enumerators in order to assist in the resolution of a given problem (the system might, for example, request that the length of the queue at a certain uninstrumented junction be measured or that a sample count be taken at a location where the automatic counter is defective). The decision to request such information, just like any decisions made by the expert system, would be based on its desire to optimise the performance of the traffic system subject to the costs of so doing being lower than the probable benefits to be derived. This again is an area where an expert system that could learn from its own experience and the constructive comments of human experts would be particularly helpful.

Air traffic control has so far attracted more attention as a possible subject for using AI techniques than has road traffic control. In the USA, proposals to develop advanced automation systems for air traffic control include new capabilities (known as AERA) for en-route detection and resolution of control problems; these will provide the controller with additional planning

information and the optional automative features, both to help in the near-term (within 20 minutes) and longer-term time horizon (Elsaesser et al, 1984).

Though the AERA proposals do not explicitly envisage the use of AI techniques, their potential is being investigated through some experimental systems. In presentations at the 1986 Meeting of the USA Transportation Research Board, possible applications of AI helpful to air traffic control were identified as: detection of collision-risk situations; suggestion of collision-avoidance manoeuvres; prevention of aircraft build-up; provision of advance information on aircraft movements; management of displays; advice on flight service procedures; checks on equipment safety; assistance in training; and detection of freak weather conditions.

A pilot example of the first of these, collision avoidance, was developed by the MITRE Corporation. Known as AIRPAC (Adviser for Intelligent Resolution of Predicted Air Conflicts), this considered conflicts between just two aircraft, and involved four successive stages: problem decomposition; tactic selection (the kernel of the system); tactic development; and manoeuvre parameter selection. The Lincoln Laboratory have developed a system for detecting wind-shear patterns from weather radar. In the UK, the Civil Aviation Authority is considering the application of IKBS techniques to flow regulation (to minimise delay), and to the use of shared air-space.

Notwithstanding these developments, it seems unlikely that expert system technology will completely replace the human air traffic controller, if only because of the suspicion that no fully automatic system will be as able as a human to cope with the unpredictable. Expert system technology is, however, likely in due course to provide valuable decision aids.

A survey control system might be linked to on-line data collection devices or might have data transmitted to it periodically from data loggers in the field. With this information it would be able to calculate distributions and variances and, knowing the desired level of precision for given data items, would then issue advice as to what extra data ought to be collected where and when. Use of the system would then ensure minimum waste of survey effort. Much of this could, of course, be achieved by conventional computing but an expert system might be better able to act on contextual information such as overall project budget, site conditions, possible trade-offs between the desired precision of different data items, marginal costs of extra data collection etc. The advice system might learn how to deal with this contextual information by receiving constructive comments from human experts on its suggested actions and by monitoring the success or otherwise of its own suggestions.

Policy Support Systems.

Expert systems can in general be regarded as decision aids, and there is a rapidly developing literature on their application in the context of decision support systems. We include here just three examples of how they might assist in a policy support role.

Multi criteria decision making involves assessing the relative value of options which are described in terms of multiple criteria - eg two alternative routings for a new road might give quite different benefits in terms of time savings, accidents, environmental effects, etc. An aid to such decision making might be provided by an expert system which, on the basis of previously logged decisions or of trade-off exercises set by the system, could devise the most appropriate way of combining the different criteria. Trade-off exercises might involve the user being asked to rank or score various packages of attributes or, more directly, to specify trade-off values (eg a 1% reduction in road accidents is of equivalent value to a 2% increase in journey time). Such a system would mimic the procedures currently recommended by professional decision analysts.

Advice on the treatment of uncertainty in forecasting or appraisal might be particularly useful. The system might calculate the net costs and benefits of a series of possible outcomes and then, having ascertained the policy maker's attitude to risk on the basis of direct questions or by deduction from his past decisions, would advise on the optimal decision. Clearly much of the benefit of such a system might be obtained from conventional computing (why has it not been done?) and it is perhaps only the ability to deduce the decision maker's attitude to risk from his past decision that is characteristic of an expert system (although even this might be achieved via more conventional techniques).

Advice on the consistency of policy might not always be welcomed by a policy maker but he might accept it more readily from a machine than from a human! A system might well be developed to investigate past policy decisions and deduce values put on various items (eg the value of commuter travel time, of human life, of clean air, of aesthetics etc). Once this has been done the system might either suggest what decision on an outstanding question would be consistent with earlier decisions or might point out inconsistencies between the proposed decision and earlier ones. The system might seek the reasons for such apparent inconsistencies and hence update its knowledge base. Thus it might say 'the proposed decision implies a much higher value on human life than any previous decision, do you mean this to be the case and if so how do you justify it?'

DISCUSSION

The successful application of AI and expert systems techniques in transport, and indeed in any applied discipline, will depend on the ease with which transport professionals adopt and use the techniques for themselves. Unfortunately, there are at present several barriers to understanding such systems. The jargon is most off-putting, there is a lot of it, several terms may be used for the same thing, and the same term may mean different things to different people. Concepts are different, and the way of thinking about a problem is different. Illustrative material often seems trivial or irrelevant.

Fortunately, a number of texts are becoming available to help demystify and explain the subject; guidance on some of these is given in the Appendix.

In course of time the techniques will probably become as much a part of the armoury of tools as statistical methods are now. However, unlike statistical methods, familiarity with the tools cannot be acquired in isolation from the computer. The ways in which the computer represents knowledge and relationships are intrinsic to the tools, and this means that their successful adoption will require much dedicated effort.

The basic concept of embodying human rules of thumb and expert judgement in a computer program is straightforward to grasp but difficult to apply. It is all too easy to fall into the trap of expecting too much of an expert system. Expert systems have outperformed humans only for very strictly defined tasks. More often they are able only to mimic the human's performance in a quite restricted part of his domain of expertise.

Many of the early expert system shells came about simply by stripping out the original knowledge base from a given expert system. For example, the shell EMYCIN came from MYCIN (here "E"="Empty"!). It is important to realise that a shell designed for one purpose may be quite inadequate for another. In particular, some of the shells treat uncertainty in a quite simplistic manner.

In some fields in civil engineering, such as construction and building, regulations provide rich scope for straightforward applications of expert systems using relatively simple shells on micros. They thus provide a good industry-wide starting point for familiarisation with what such systems can do. Similarly simple transport applications, on the other hand, are rather restrictive in their appeal and the transport applications we have identified above are likely to be rather demanding of expert system technology. In this respect we do not accept d'Agapeyeff's (1984) conclusion that "simpler expert systems are practical" for most applications of interest in the transport sector.

Serious transport applications will place several demands on the developers of expert system technology. Foremost, given the importance of number crunching in most transport applications, must be the development of systems or languages that are much better able to integrate the handling of numbers with rules and relationships. Secondly, since systems fall into disuse if cruder but quicker techniques will suffice, expert system technology needs to be speeded-up; not only for real-time applications such as traffic monitoring and control, but also more generally. But perhaps the greatest challenge facing the developers of expert system technology is to improve the systems' ability to learn from their experience in being used. Several times in our list of transport applications we have suggested that this ability is particularly beneficial; but it is a strange anomaly of the artificial intelligence field that, though simple applications of such 'automated learning' processes can be found in most introductory AI texts and even programmed in Basic for home micros, the introduction of automated learning techniques in serious applications is extremely difficult, and generally not provided for in current software tools.

Alongside the need to develop expert system technology itself, is the need to develop the man-machine interface via enhanced dialogue facilities. It is similarly necessary to develop the skills and procedures for acquiring the knowledge that is to be

built into the system. This need lies at the heart of any expert system, but surprisingly little progress has been made in the development of systematic knowledge acquisition methodologies (for reviews on which see Welbank, 1983, and Bourne and Sztipanovits, 1985). One implication of this is that the development of an expert system should not be left solely to an expert system specialist; general knowledge of the subject will be important in identifying the right sort of questions to ask, of whom and in what order.

Expert system technologies provide the tools to enable a wider range of problems to be tackled (or to enable a given problem to be tackled more comprehensively). The goal must not be to produce "expert systems" per se; but rather to produce systems in which expert system tools are used for what they are good at, and algorithmic procedures used for what they are good at. Arguments as to whether one should use 'shells' or 'languages', or whether the eventual program is worthy to be called an expert system or not, will eventually be seen to be rather sterile. Expert system technology will in the end only be thought useful if embedded in an approach that confronts the whole of a problem, not just part of it.

Whatever the end product be called, the need remains for practitioners and clients to become better acquainted with what expert system methodology can do, and for their analysts to become well versed not only in what it can do, but how different tools contribute to doing it. The development of pilot, but not trivial, demonstration projects should contribute significantly to this end.

APPENDIX: SOME INTRODUCTORY SOURCES

The need to de-mystify and explain the subject of artificial intelligence and expert systems is now being addressed at several levels. Home computer users have become greatly interested in the techniques, partly perhaps because some of the AI techniques (such as text compression and search procedures) are exploited in the advanced adventure games now available. This has resulted in a variety of easy-to-use introductory sources on the subject, covering not only books (eg Naylor, 1983) but also viewdata-based user groups (MICRONET, 1986) and useful, though somewhat limited, implementations of LISP and micro-PROLOG on home micros.

At a more serious level, there are a number of introductory articles in journals for specific application areas; see for example our own report (Wheatley, 1984) for transport, and Hinde (1985) for Operational Research. Several good text books are available describing the essential background to and nature of the subject: see for example Rich (1983) and Bonnet (1986). For those wanting to get the feel for commercial expert systems, some 'starter packs' are now available, ranging from those that cope only with production rules to those that incorporate scaled-down versions of commercial systems; that issued for the Alvey Directorate (1984) by the National Computing Centre contains what Wigan (1985) describes as 'crippled' versions of Expert-Ease, ESP Adviser and Micro-Expert, and Micro

Synics. Useful reviews of many of the leading expert system shells can be found in Allwood et al (1985) and Wigan (1985); summary details can also be obtained from the Institute on request. A useful catalogue of a wider range of artificial intelligence tools is given by Bundy (1986).

REFERENCES

- d'AGAPEYEFF, A. (1984) Report to the Alvey Directorate on a short survey of expert systems in UK business. Alvey News, Supplement to Issue No. 4.
- ALLWOOD, R.J., STEWART, D.J., HINDE, C. and NEGUS, B. (1985) Evaluation of expert system shells for construction industry applications. Department of Civil Engineering, University of Loughborough.
- ALVEY DIRECTORATE (1984) IKBS Starter Pack. National Computing Centre, Manchester.
- BOBROW, D.G., KAPLAN, R.M., KAY, M., NORMAN, D.A., THOMPSON, H. and WINOGRAD, T. (1977) GUS, a frame-driven dialog system. Artificial Intelligence, 8, 157.
- BONNET, A. (1986) Artificial intelligence: Promise and performance. Prentice Hall International, London. ISBN 0-13-048869-0, £9.95.
- BOURNE, J.R. and SZTIPANOVITS, J. (1985) Strategies for knowledge representation, manipulation and acquisition in expert systems. Seventh Ann. Conf. of the Engineering in Medicine and Biology Society, Institute of Electrical and Electronic Engineering, 1165-1169.
- BOYCE, D.E. (ed) (1985) Transportation research: the state of the art and research opportunities. Transportation Research, 19A(5/6), Special Issue.
- BUNDY, (1986) Catalogue of artificial intelligence tools. Springer. ISBN 3-54-013938-9, £13.25.
- CENTRE FOR EDUCATIONAL SOCIOLOGY (1984) QUESTMAST users' manual. Centre for Educational Sociology, University of Edinburgh.
- DEPARTMENT OF INDUSTRY (1982) A programme for advanced information technology: The Report of the Alvey Committee. HMSO, London, ISBN 0-11-513653-3, £4.30.
- ELSAESSER, C., GISCH, A.H., HAINES, A.L. and SWEDISH, W.J. (1984) Description of AERA 1 capabilities. MTR-84W88, Metrek Division, Mitre Corporation, McLean, Virginia, USA (unpublished).
- FEIGENBAUM, E.A. and MCCORDUCK, P. (1984) The fifth generation: Artificial intelligence and Japan's computer challenge to the world. Pan Books, London.
- HARRIS, R.A., COHN, L.F. and BOWLBY, W. (1985) An application of artificial intelligence in highway noise analysis. Transportation Research Record 1033.
- HENDRICKSON, C.T., REHAK, D.R. and FENVES, S.J. (1985) Expert systems in transportation systems engineering. Research Report, Department of Civil Engineering, Carnegie-Mellon University.
- HINDE, C.J. (1985) Artificial intelligence and expert systems. In: Rand, G.K. and Eglese, R.W. (eds) Further developments in Operational Research. Pergamon Press, Oxford.
- KWAN, R.S.K. (1985) Approaches to expert systems for passenger enquiries using SAGE. Technical Note 163, Institute for Transport Studies, University of Leeds, Leeds.
- LOGIE, M. and NEFFENDORF, H. (1984) Expert systems in transportation. PTRC Summer Annual Meeting, Brighton 1984.
- MACLEAN, I. (1983) Proceedings of the Transport Statistics Conference, November 1983. Annual Conference of the Statistics Users Council. IMAC Research, London.

- MANDANI, A.; EFSTATHIOU, J. and PANG, D. (1985) Inference under uncertainty. Unpublished report, Queen Mary College, University of London.
- MICRONET (1986) AI Queries. PRESTEL, frame *8006055 et seq.
- MOSTOW, J. (1985) Towards better models of the design process. The AI Magazine, Spring 1985, 44-56.
- NAYLOR, C. (1983) Build your own expert system. Sigma, Wilmslow, ISBN 0-905104-41-2, £6.95.
- O'KEEFE, R.M., BELTON, V. and BALL, T. (1985) Getting into expert systems. QSS Discussion Paper No. 71, Board of Studies in Quantitative Social Science and Management Science, University of Kent at Canterbury.
- POLAK, J.W. and JESSOP, A. (1986) Towards expert systems for goods distribution management. Conference Universities Transport Study Group, held at UWIST, Cardiff, Jan. 1986 (unpublished).
- RICH, E. (1983) Artificial intelligence. McGraw-Hill.
- SOFTWARE SCIENCES Ltd, WOOTTON JEFFREYS plc, and INSTITUTE FOR TRANSPORT STUDIES, UNIVERSITY OF LEEDS (1985) Expert enquiry systems. A proposal to the Alvey Directorate to form a Community Club in Transport. Software Sciences Ltd, Farnborough.
- WELBANK, M. (1983) A review of knowledge acquisition techniques for expert systems. British Telecom Research Laboratory, Martlesham Consultancy Services, Ipswich.
- WELCH, J.G. and BISWAS, M. (1986) Application of expert systems in the design of bridges. 65th Annual Meeting, Transportation Research Board, Washington DC, January 1986.
- WHEATLEY, M.D. (1984) Expert systems in transport. Part 1: An introduction to expert systems. Working Paper 178, Institute for Transport Studies, University of Leeds, Leeds.
- WIGAN, M.R. (1983) Information technology and transport: what research needs to be started now? Working Paper 172, Institute for Transport Studies, University of Leeds, Leeds.
- WIGAN, M.R. (1985) Knowledge based systems tools on microsystems. Research Report ARR 134, Australian Road Research Board, Vermont, Victoria, Australia (43pp). ISBN 0-86910-206-0 (Report); ISBN 0-86910-207-9 (Microfiche).
- WILLIAMSON, R.H. and MILLER, A.J. (1981) A personalised public transport directory. Annual Public Transport Symposium, University of Newcastle upon Tyne, July 1981.
- WOOD, P. (1985) A knowledge based system for transit bus maintenance. Transportation Research Record 1019, 85-91.
- WOOTTON, H.J. and BRETT, A.C. (1986) Route information systems - signposts of the future? Proceedings of the Second International Conference on Road Traffic Control, Institution of Electrical Engineers, April 1986, London.
- YEH, C-I., RITCHIE, S.G., SCHNEIDER, J.B. and NIHAN, N.L. (1985) Knowledge-based expert systems in transportation planning and engineering: an initial assessment. Annual Meeting, Transportation Research Board, Washington D.C., January 1985.
- ZOZAYA-GOROSTIZA, C. and HENDRICKSON, C. (1986) An expert system for traffic signal assistance. Research Report, Carnegie-Mellon University.

AI/PAP (11)
25.4.86

APPENDIX A LIMITATIONS OF SOME EXPERT SYSTEMS SHELLS

This review of some limitations of expert systems shells refers to those that were available in 1984 and 1985.

1. Too much traded-off

The desire for simplicity and low cost have produced systems that are too inflexible and lacking in generality to be of practical use in contexts for which they were not specifically designed.

2. Shallow knowledge representation

Simple rule-based systems (eg SAVOIR, SAGE) have insufficient expressive power for knowledge about objects that have complicated structures and relationships. For example, whilst it is possible to express inherited characteristics via simple rules, this would require a lot of rules, and make processing slow. More appropriate methods of representation would be with frame-based or object-based systems.

The inductive system EXERT-EASE requires that the data conforms to deterministic rather than stochastic rules. This inability to allow for data error or model error in inferring rules is a fundamental drawback. Even without that problem, the validity of the decision trees produced is questionable, and the inability to modify them undesirable.

3. Rigid inference control

Standard forward/backward chaining techniques are not sufficient. In backward chaining, for example, one typically has to say which goal is to be evaluated first; in forward chaining one may want to branch out.

Some shells provide extra control facilities over the flow of investigations. However, these extra control structures are generally too rigid, being able to cope with limited problems, but hopeless for others. For example, one may want to change goals in mid-stream (eg, to arrive by 1500 on a train with a restaurant car, instead of by 1200): in SAGE (for example) one would be forced to start again.

4. Inadequate interactive support environments

The fixed format for questions and user prompts is often confusing and easily mis-interpreted by the user. For example, displaying a statement and asking 'is this true?' is often not a very natural way of replying, especially if numbers are involved. Entering answers on part of the screen remote from the question may also be

confusing. (For example, for a railway track maintenance advisory system, the US Army's Construction Engineering Research Laboratory found it necessary to replace the 'front-end' to ESP Adviser by one that was easier to use).

Most responses can only be accepted in simple terms, like 'Yes/No', a single keyword, integer numbers, menu selection. Even tailoring the system to one's own preferred simple method is usually impossible.

Rather than these simple methods, it is often more natural to want to provide several pieces of related information at once; this implies the ability to enter several keywords, perhaps with numbers as well (eg Leave by 11am, 1st class).

5. Insufficient explanation facilities

Two types of explanation facility are needed, but often only one is provided. These are "why I am asking you this" type, and "how I reached this conclusion" type. Only the latter needs to list the rules that have been triggered; and such lists are often difficult to comprehend. Ways of summarising these rules clearly and concisely are needed.

6. Primitive interfaces to other software

Interfaces to two kinds of software are needed: to the language in which the expert system is written, so that it can be enhanced; and to other, more conventional languages, usually for numeric calculations.

The following limitations are apparent:

- no interfaces at all in some shells
- awkward mechanisms for establishing the linkages and invoking the external modules (eg in SAVOIR one has to invoke the interrupt mechanism)
- restrictions on the data types that can be passed between the shell and the external modules (eg most pass only integers; SAGE, though it can link to a Pascal program, can only pass real numbers to it)
- restrictions on the ability to use the numbers passed from the external program (eg SAGE has to use the number immediately; it cannot be stored).

7. Conclusion

The main conclusion to be drawn is that the operational requirements for an expert system field to be well specified, and the characteristics of the shells available to be well described, to ensure that a shell is suitable for a given task. The early shells do not seem well suited to operational tasks.